

Forward KL Regularized Preference Optimization for Aligning Diffusion Policies

Zhao Shan^{1,2*}, Chenyou Fan^{1,3*}, Shuang Qiu⁴, Jiyuan Shi¹, Chenjia Bai^{1,5†}

¹Institute of Artificial Intelligence (TeleAI), China Telecom

²Tsinghua University

³Northwestern Polytechnical University Xi'an

⁴Hong Kong University of Science and Technology

⁵ShenZhen Research Institute of Northwestern Polytechnical University

Abstract

Diffusion models have achieved remarkable success in sequential decision-making by leveraging the highly expressive model capabilities in policy learning. A central problem for learning diffusion policies is to align the policy output with human intents in various tasks. To achieve this, previous methods conduct return-conditioned policy generation or Reinforcement Learning (RL)-based policy optimization, while they both rely on pre-defined reward functions. In this work, we propose a novel framework, Forward KL regularized Preference optimization for aligning Diffusion policies, to align the diffusion policy with preferences directly. We first train a diffusion policy from the offline dataset without considering the preference, and then align the policy to the preference data via direct preference optimization. During the alignment phase, we formulate direct preference learning in a diffusion policy, where the forward KL regularization is employed in preference optimization to avoid generating out-of-distribution actions. We conduct extensive experiments for MetaWorld manipulation and D4RL tasks. The results show our method exhibits superior alignment with preferences and outperforms previous state-of-the-art algorithms.

Introduction

In solving sequential decision-making problems, Reinforcement Learning (RL) algorithms typically adopt Gaussian or deterministic policy classes in policy optimization. Although such a policy class has successfully solved various challenging tasks (Fujimoto, Meger, and Precup 2019; Lange, Gabel, and Riedmiller 2012), it can be limited in learning multi-modal policies and lead to sub-optimal behaviors in complex environments (Janner et al. 2022). Recently, diffusion models (Ho, Jain, and Abbeel 2020; He et al. 2024) demonstrated superior performance compared to previous policy classes, especially for offline RL where the dataset is collected by a mixture of policies (Fu et al. 2020; Qin et al. 2022; Yuan et al. 2024a; Fan et al. 2024), or the embodied manipulation tasks where the policy should imitate diverse human behaviors (Chi et al. 2023; Pearce et al. 2023). Leveraging the highly expressive model capabilities,

the diffusion policy exhibits strong capabilities in modeling complex behaviors.

A central problem in learning a diffusion policy is aligning the policy output with human intents (Shin, Dragan, and Brown 2023; Yuan et al. 2024b; Yu et al. 2024). Existing methods can be roughly divided into three categories. (i) Several methods collect expert trajectories and directly learn a diffusion policy from the expert dataset (Reuss et al. 2023; Ze et al. 2024), while they usually require a large number of expert trajectories to model human behaviors. To ensure the diffusion policy generates actions with desired properties, (ii) several methods adopt conditional generation by employing the cumulative return as the condition in modeling the trajectory (Yuan et al. 2023; Ajay et al. 2023), and actions with high returns can be generated via guided sampling (Ho and Salimans 2021; Nichol and Dhariwal 2021). (iii) Other works consider diffusion policy as an actor in an RL framework and perform policy improvement by maximizing cumulative rewards. Then actions can be sampled with high Q -values (Wang, Hunt, and Zhou 2023; Kang et al. 2023). Although the latter two methods can learn policies from an arbitrary dataset, they rely on manually defined reward functions for each task, which can be difficult to obtain in embodied tasks (Dong et al. 2024; Gu et al. 2023). Motivated by recent advances in Reinforcement Learning from Human Feedback (RLHF) (Ouyang et al. 2022; Bai et al. 2022) that align human preference with Large Language Models (LLMs) (OpenAI 2023; Touvron et al. 2023), we wonder whether the diffusion policy can be aligned with human intent directly from a collection of preference data, which removes the potentially incorrect assumption that the reward function alone drives human preferences (Knox et al. 2024; Lambert et al. 2024) and aims to guide the diffusion policy from preferences directly.

In this work, we propose a novel framework, named Forward KL regularized Preference optimization for aligning Diffusion policies (**FKPD**), to align the diffusion policy with preferences. To achieve this, FKPD learns a basic diffusion policy from the offline dataset, then aligns the policy to preferences via policy optimization with a forward KL regularization with respect to the offline dataset. In the first stage, the basic policy recovers the complex action distribution in the dataset without considering the preference la-

*These authors contributed equally.

†Corresponding author

bel, which can be multi-modal to capture the fidelity of behaviors. In the second stage, we align the basic diffusion policy to the preference data through Direct Preference Optimization (DPO) in a maximum entropy RL framework, which directly updates the policy output to match the preference model. The two stages in FKPD are complementary since the first stage learns the full distribution of actions, and the second stage searches policy based on the full distribution and converges to more narrowly distributed actions that align well with preferences. The basic policy learned in the first stage provides multi-modal policy initialization and a foundation for approximating the reverse process using the forward process for preference optimization.

The remaining challenge of this work lies in formulating the objective function for the alignment phase. To address the intractable likelihoods of diffusion policies, we approximate the optimization objective by sampling from the forward chain of the diffusion model to derive a tractable objective for preference optimization. In addition, we employ the forward KL regularization, which is more efficient for a diffusion policy than the reverse KL regularization adopted in most existing DPO literature (Rafailov et al. 2024b; Wallace et al. 2023), to constrain the distance between the diffusion policy and the behavior policy, thereby addressing the out-of-distribution (OOD) issues in the alignment phase.

Our contribution can be summarized as follows. (i) We provide a novel framework to learn diffusion policies directly from preference data without relying on rewards and demonstrations. Specifically, we first propose a two-stage process by first learning a basic policy and then performing preference optimization on it. (ii) We propose the DPO objective for diffusion policies in the alignment phase which is shown to be more efficient than most existing methods (Rafailov et al. 2024b; Wallace et al. 2023). Specifically, we employ forward KL regularization in the alignment of diffusion policies, which is shown to be more effective in addressing OOD issues. (iii) We conduct experiments on MetaWorld (Yu et al. 2020) and D4RL (Fu et al. 2020) tasks, and the results show that FKPD exhibits superior alignment with preferences and outperforms previous state-of-the-art methods.

Preliminaries

Preference-based Reinforcement Learning

We consider the general Preference-based Reinforcement Learning (PBRL) problem within a reward-free Markov Decision Process (MDP) $\mathcal{M}/r = (\mathcal{S}, \mathcal{A}, p, \gamma)$ with state space \mathcal{S} , action space \mathcal{A} , transition dynamics $p(s^{t+1}|s^t, a^t)$ and discount factor γ . The goal is to learn a policy $\pi(a|s)$ that maximizes an expert user’s reward. We are provided with an offline dataset \mathcal{D} resembling the setting of offline reinforcement learning, except that the dataset does not contain reward information. Instead, we are also given a preference dataset $\mathcal{D}_{\text{pref}} = \{(\sigma_i^+, \sigma_i^-)\}_{i=1}^n$ consisting of pairs of segments sorted according to user preferences. Here $\sigma = (s^1, a^1, s^2, a^2, \dots, s^k, a^k)$ is a length- k segment sampled from \mathcal{D} . We use $\sigma^+ \succ \sigma^-$ to indicate that σ^+ is preferred to σ^- . In the following sections, σ^+ and σ^- are re-

ferred to as “winning” and “losing” segments, respectively.

Diffusion Model for Decision-Making

Diffusion-based generative models (Ho, Jain, and Abbeel 2020) assume $p_\theta(\mathbf{x}_0) := \int p_\theta(\mathbf{x}_{0:T})d\mathbf{x}_{1:T}$, where $\mathbf{x}_1, \dots, \mathbf{x}_T$ are latent variables of the same dimensionality as the data $\mathbf{x}_0 \sim p(\mathbf{x}_0)$. A forward diffusion chain gradually adds noise to the data $\mathbf{x}_0 \sim q(\mathbf{x}_0)$ in T steps with a pre-defined variance schedule β_t , expressed as $q(\mathbf{x}_{1:T}|\mathbf{x}_0) := \prod_{t=1}^T q(\mathbf{x}_t|\mathbf{x}_{t-1})$, where $q(\mathbf{x}_t|\mathbf{x}_{t-1}) := \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I})$.

A reverse chain, constructed as $p_\theta(\mathbf{x}_{0:T}) := \mathcal{N}(\mathbf{x}_T; \mathbf{0}, \mathbf{I}) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$, is then optimized by maximizing the evidence lower bound (ELBO) defined as $\mathbb{E}_q \left[\ln \frac{p_\theta(\mathbf{x}_0)}{q(\mathbf{x}_{1:T}|\mathbf{x}_0)} \right]$. After training, sampling from the diffusion model consists of sampling from \mathbf{x}_T and running the reverse diffusion chain from $t = T$ to $t = 0$. Diffusion models can be extended to conditional generative models by conditioning $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, c)$. Since there are two different types of timesteps in this work, one for the diffusion process and one for the MDP, we use superscripts to denote trajectory timesteps and subscripts to denote diffusion timesteps. We denote the diffusion-base policy via the reverse process of a conditional diffusion model as $\pi_\theta(a|s) = \pi_\theta(a_{0:T}|s) = \mathcal{N}(a_T; \mathbf{0}, \mathbf{I}) \prod_{t=1}^T \pi_\theta(a_{t-1}|a_t, s)$, where the end sample of the reverse chain, a_0 , serves as the action for RL tasks. When we use a diffusion policy to fit the behavior policy on \mathcal{D} (behavior clone), we adopt the simplified objective proposed by Ho, Jain, and Abbeel (2020)

$$\mathcal{L}_{\text{BC}} = \mathbb{E}_{t \sim \mathcal{U}_1, (s, a) \sim \mathcal{D}} \left[\left\| \epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t}a + \sqrt{1 - \bar{\alpha}_t}\epsilon, s, t) \right\|^2 \right], \quad (1)$$

where $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\bar{\alpha}_t$ is a scale factor defined in Ho, Jain, and Abbeel (2020).

Method

The overall scheme of our method is illustrated in Fig. 1. We first pre-train a basic diffusion policy via behavior cloning from the preference-free dataset \mathcal{D} , followed by aligning this policy with the preference dataset $\mathcal{D}_{\text{pref}}$. For the behavior clone phase, we train a diffusion policy π_θ by solving $\max_{(s, a) \in \mathcal{D}} [\log \pi_\theta(a|s)]$. This is a standard process of training a diffusion policy that fits the behavior policy without reward or preference conditions (Wang, Hunt, and Zhou 2023) and will not be elaborated further. In terms of the alignment of a diffusion policy, we utilize the DPO framework, eliminating the necessity for an explicit reward model extracted from the preference dataset, thus bypassing the risk associated with an inaccurate reward model. In this section, we first give the objective of the alignment phase, referred to as $\mathcal{L}_{\text{DPO-FK}}$, for a general generative policy from the perspective of maximum entropy RL, where we introduce a forward KL regularization to avoid the generation of OOD actions. Subsequently, we give a practical approximation of

$\mathcal{L}_{\text{DPO-FK}}$ for a diffusion policy. Moreover, we provide an intuitive explanation of why forward KL regularization is superior to another common regularization method, namely, reverse KL regularization (Wallace et al. 2023).

DPO with Forward KL Regularization

For a general generative policy, a length- k segment $\sigma = (s^1, a^1, s^2, a^2, \dots, s^k, a^k)$ can be regarded as generated from a conditional generative model, where (a^1, a^2, \dots, a^k) acts as the sample and (s^1, s^2, \dots, s^k) as the condition. Here, the generative model is a stack of identical generative policies $\pi_\theta(a|s)$. In order to generate a segment of actions, we parallelize a stack of individual basic generative policies. Without ambiguity, we still denote the generative model producing a segment σ as the generative policy π_θ . Moreover, we omit (s^1, s^2, \dots, s^k) in a segment and denote $\pi_\theta(\sigma) = \prod_{i=1}^k \pi_\theta(a^i|s^i)$ for brevity.

During the alignment phase, we aim to optimize π_θ such that the preference score of a generated segment σ is maximized. At the same time, we also need to control π_θ to prevent it from deviating too far from the behavior policy, in order to avoid OOD issues. To this end, we use the following objective during the alignment phase:

$$\max_{\pi_\theta} \underbrace{\mathbb{E}_{\sigma \sim \pi_\theta} [r(\sigma)]}_{\text{reward maximization}} - \mu \underbrace{\mathbb{D}_{\text{KL}}[\pi_{\text{ref}} \parallel \pi_\theta]}_{\text{forward KL}}, \quad (2)$$

where $r(\cdot)$ is a latent score function modeling human preference, π_{ref} is the policy obtained from the behavior cloning phase serving as a precise approximation of the behavior policy and μ is a balance factor. Here we use a forward KL regularization to control π_θ during the alignment phase. However, most existing work (Wallace et al. 2023; Rafailov et al. 2024b) employs a reverse KL regularization, where they use an objective as

$$\max_{\pi_\theta} \underbrace{\mathbb{E}_{\sigma \sim \pi_\theta} [r(\sigma)]}_{\text{reward maximization}} - \mu \underbrace{\mathbb{D}_{\text{KL}}[\pi_\theta \parallel \pi_{\text{ref}}]}_{\text{reverse KL}}. \quad (3)$$

Recent studies indicate that reverse KL and forward KL regularization possess distinct characteristics, with reverse KL exhibiting mode-seeking behavior and forward KL demonstrating mass-covering tendencies (Sun and van der Schaar 2024). Furthermore, forward KL has been shown to better balance alignment performance and generation diversity (Wang et al. 2024). In this paper, we will reveal another significant advantage of forward KL regularization in the alignment of diffusion policies: its superior effectiveness in avoiding OOD issues, which will be explained in subsequent sections.

To optimize Eq. (2) directly using preference data, we handle two terms in Eq. (2) separately. For the forward KL regularization term, we simplify it as

$$\begin{aligned} & \mathbb{D}_{\text{KL}}[\pi_{\text{ref}} \parallel \pi_\theta] \\ &= \mathbb{E}_{\sigma \sim \pi_{\text{ref}}} [\log \pi_{\text{ref}}(\sigma)] - \mathbb{E}_{\sigma \sim \pi_{\text{ref}}} [\log \pi_\theta(\sigma)] \\ &\approx \underbrace{\mathbb{E}_{\sigma \sim \pi_{\text{ref}}} [\log \pi_{\text{ref}}(\sigma)]}_{\text{constant}} - \mathbb{E}_{\sigma \sim \mathcal{D}} [\log \pi_\theta(\sigma)]. \end{aligned} \quad (4)$$

The forward KL regularization is divided into the entropy of π_{ref} , which is a constant, and the cross entropy of π_{ref} and π_θ . To compute the cross-entropy term, we can sample σ from \mathcal{D} instead of from π_{ref} , as π_{ref} serves as a precious approximation of behavior policy.

In terms of the reward maximization term in Eq. (2), we consider π_θ from the perspective of maximum entropy RL (Haarnoja et al. 2018), leading to the following problem:

$$\max_{\pi_\theta} \mathbb{E}_{\sigma \sim \pi_\theta} [r(\sigma) - \rho \log \pi_\theta(\sigma)], \quad (5)$$

where $r(\cdot)$ is a latent score function modeling human preference, and ρ is a temperature parameter. This problem has a closed-form solution (see Appendix 1.1 for details), given by

$$\pi_\theta^*(\sigma) = \frac{1}{Z} \exp\left(\frac{r(\sigma)}{\rho}\right), \quad (6)$$

where Z is a normalization constant. In other words, the unknown score function r can be expressed by the optimal generative policy as $r(\sigma) = \rho \log(\pi_\theta^*(\sigma)) + \rho \log(Z)$. According to Bradley-Terry model (Bradley and Terry 1952), the human preference model on $\mathcal{D}_{\text{pref}}$ is written as

$$\begin{aligned} P(\sigma^+ \succ \sigma^-) &= \text{Sigmoid}(r(\sigma^+) - r(\sigma^-)) \\ &= \text{Sigmoid}(\rho \log \pi_\theta^*(\sigma^+) - \rho \log \pi_\theta^*(\sigma^-)). \end{aligned} \quad (7)$$

Then we can derive the DPO objective for a maximum entropy generative policy employing the principle of contrastive learning, i.e., minimize the cross-entropy loss of the preference model Eq. (7) on $\mathcal{D}_{\text{pref}}$, which is written as

$$\begin{aligned} \mathcal{L}_{\text{DPO}}(\pi_\theta, \mathcal{D}_{\text{pref}}) &= \mathbb{E}_{(\sigma^+, \sigma^-) \sim \mathcal{D}_{\text{pref}}} \\ &- \text{Logsigmoid}(\rho (\log \pi_\theta(\sigma^+) - \log \pi_\theta(\sigma^-))). \end{aligned} \quad (8)$$

In the following sections, we denote $\text{Logsigmoid}(\rho(\cdot))$ as $F(\cdot)$ for the sake of brevity. Next, we combine the above derivations together and obtain the following objective for the alignment phase:

$$\begin{aligned} \mathcal{L}_{\text{DPO-FK}}(\pi_\theta, \mathcal{D}_{\text{pref}}) &= \\ \mathcal{L}_{\text{DPO}}(\pi_\theta, \mathcal{D}_{\text{pref}}) - \mu \mathbb{E}_{\sigma \in \mathcal{D}} \log \pi_\theta(\sigma). \end{aligned} \quad (9)$$

Practical Approximation for Diffusion Policy

It is fairly hard to calculate Eq. (9) for a diffusion policy. We will next give a practical approximation. In fact, we only need to tackle the first term since the regularization term is bounded by its ELBO (Ho, Jain, and Abbeel 2020). However, the first term cannot be simplified by the same technique since the signs of the first likelihood term and the second one in Eq. (8) are opposite. To tackle this issue, we introduce all latent variables $\sigma_1, \sigma_2, \dots, \sigma_T$ in the reverse process similar to Wallace et al. (2023). By defining the score function on the whole reverse chain as $R(\sigma_{0:T})$, then the original score function r can be calculated by marginalizing out all latent variables as

$$r(\sigma_0) = \mathbb{E}_{\sigma_{1:T} \sim \pi_\theta(\sigma_{1:T}|\sigma_0)} [R(\sigma_{0:T})]. \quad (10)$$

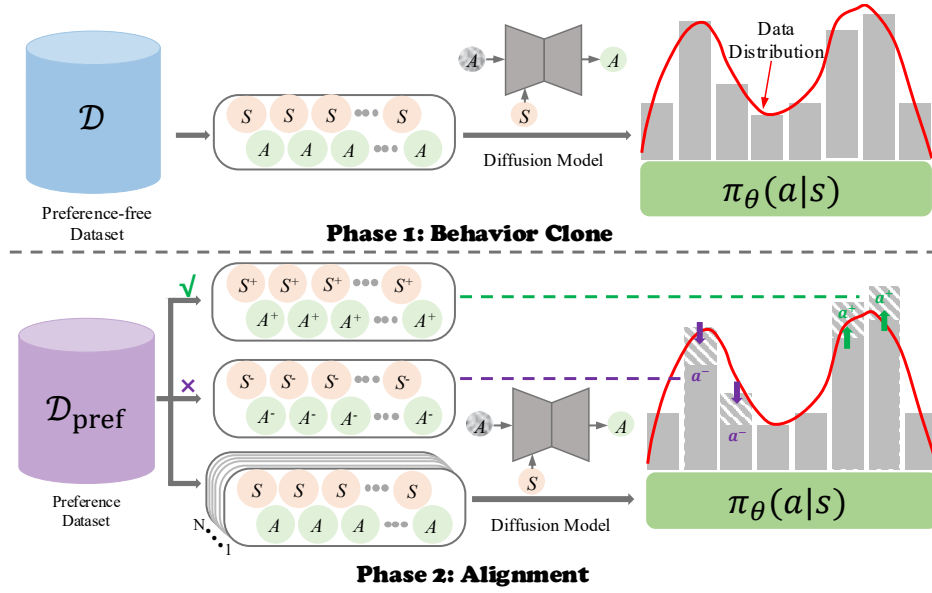


Figure 1: Overall scheme of FKPD: FKPD comprises two phases: the behavior cloning phase and the alignment phase. In the behavior clone phase, FKPD fits a diffusion policy to a preference-free dataset \mathcal{D} . Subsequently, in the alignment phase, FKPD employs a direct preference optimization method to align the diffusion policy with a preference dataset $\mathcal{D}_{\text{pref}}$. Throughout this process, the diffusion policy is required to maintain a forward KL distance with respect to the distribution of $\mathcal{D}_{\text{pref}}$

Then we perform the same process as in Eqs. (5)-(8), yielding the DPO objective for a diffusion policy as (see Appendix 1.2 for details),

$$\mathcal{L}_{\text{DPOdiff}}(\pi_\theta, \mathcal{D}_{\text{pref}}) = \mathbb{E}_{\mathcal{D}_{\text{pref}}} \left[-F \left(\mathbb{E}_{\sigma_{1:T}^+ \sim \pi_\theta(\sigma_{1:T}^+ | \sigma_0^+)} \log \pi_\theta(\sigma_{0:T}^+) - \mathbb{E}_{\sigma_{1:T}^- \sim \pi_\theta(\sigma_{1:T}^- | \sigma_0^-)} \log \pi_\theta(\sigma_{0:T}^-) \right) \right]. \quad (11)$$

Now, it is also challenging to optimize Eq. (11) directly, as it needs to take an expectation over the whole reverse process, which is hard to track. To tackle this issue, we use the forward process $q(\sigma_{1:T} | \sigma_0)$ to approximate the reverse process, yielding the following objective:

$$\mathcal{L}_{\text{DPOdiff}}(\pi_\theta, \mathcal{D}_{\text{pref}}) \approx \mathcal{L}_1(\pi_\theta, \mathcal{D}_{\text{pref}}) \triangleq \mathbb{E}_{(\sigma^+, \sigma^-) \sim \mathcal{D}_{\text{pref}}} \left[-F \left(\mathbb{E}_{\sigma_{1:T}^+ \sim q(\sigma_{1:T}^+ | \sigma_0^+)} \log \pi_\theta(\sigma_{0:T}^+) - \mathbb{E}_{\sigma_{1:T}^- \sim q(\sigma_{1:T}^- | \sigma_0^-)} \log \pi_\theta(\sigma_{0:T}^-) \right) \right]. \quad (12)$$

This approximation makes sense for the following two considerations: on the one hand, π_θ has been trained to convergence on \mathcal{D} , indicating a close approximation between the backward and forward processes; on the other hand, the forward KL regularization constrains π_θ from deviating excessively from the pre-trained policy during the alignment phase.

For further simplification, we decompose the reverse process in Eq. (12) as $\pi_\theta(\sigma) = \pi_\theta(\sigma_T) \prod_{t=1}^T \pi_\theta(\sigma_{t-1} | \sigma_t)$, and utilize Jensen's inequality and the convexity of Logsigmoid to take the expectation over σ_t^+ and σ_t^- out of Logsigmoid, and employ Gaussian reparameterization for further

simplification, yielding the following bound of Eq. (12) (Appendix 1.3), as

$$\mathcal{L}_1(\pi_\theta, \mathcal{D}_{\text{pref}}) \leq \mathcal{L}_2(\pi_\theta, \mathcal{D}_{\text{pref}}) \triangleq -\mathbb{E}_{(\sigma^+, \sigma^-) \sim \mathcal{D}_{\text{pref}}} \mathbb{E}_{t \sim \mathcal{U}(1, T), \epsilon_t^+, \epsilon_t^- \sim \mathcal{N}(0, I)} F \left(-T \left(\|\epsilon_t^+ - \epsilon_\theta(\sigma_t^+, t)\|^2 - \|\epsilon_t^- - \epsilon_\theta(\sigma_t^-, t)\|^2 \right) \right), \quad (13)$$

where $\sigma_t = \sqrt{\bar{\alpha}_t} \sigma_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$. It is obvious that the constant T can be factored into ρ . We next substitute Eq. (13) into Eq. (9) and bound the regularization term in Eq. (9) using its ELBO, resulting in the complete objective:

$$\mathcal{L}_{\text{DPOdiffFK}}(\pi_\theta, \mathcal{D}_{\text{pref}}) = \mathcal{L}_2(\pi_\theta, \mathcal{D}_{\text{pref}}) + \mu \mathbb{E}_{\sigma \in \mathcal{D}} \mathbb{E}_{t \sim \mathcal{U}(1, T), \epsilon_t \sim \mathcal{N}(0, I)} \left(\|\epsilon_t - \epsilon_\theta(\sigma_t, t)\|^2 \right). \quad (14)$$

Finally, we move the second term of Eq. (14) into Logsigmoid, leading to our final loss for the alignment phase:

$$\mathcal{L}_{\text{FKPD}} = -\mathbb{E}_{(\sigma^+, \sigma^-) \sim \mathcal{D}_{\text{pref}}} \mathbb{E}_{t \sim \mathcal{U}(1, T), \epsilon_t^+, \epsilon_t^- \sim \mathcal{N}(0, I)} \text{Sigmoid} \left(-\rho \left(\underbrace{\|\epsilon_t^+ - \epsilon_\theta(\sigma_t^+, t)\|^2 - \|\epsilon_t^- - \epsilon_\theta(\sigma_t^-, t)\|^2}_{\text{preference}} + \underbrace{\mu \mathbb{E}_{\sigma \in \mathcal{D}} \mathbb{E}_{\epsilon_t \sim \mathcal{N}(0, I)} \|\epsilon_t - \epsilon_\theta(\sigma_t, t)\|^2 - b}_{\text{regularization}} \right) \right), \quad (15)$$

where b is a constant bias term to ensure the Sigmoid function operates within its linear regime. It should be noted that

we perform the final merge solely to facilitate the finding of the hyperparameter μ since the preference term and the regularization term operate at the same scale. In practice, the regularization term can be constrained in a small regime around zero, which will not compromise the accuracy of the Bradley-Terry preference model.

Now we can provide an intuitive explanation of how $\mathcal{L}_{\text{FKDP}}$ aligns the pretrained behavior clone policy with $\mathcal{D}_{\text{pref}}$. We first define $\mathbb{E}_{\sigma,t} \|\epsilon_t - \epsilon_\theta(\sigma_t, t)\|^2$ as the Denoising Mean Square Error, abbreviated as **D-MSE**. A smaller D-MSE for a segment implies a higher likelihood that the diffusion policy will generate that segment. There are two terms inside the Logsigmoid function: the preference term describes the discrepancy in the D-MSE for the winning segment relative to the losing segment, and the regularization term represents the average D-MSE for dataset \mathcal{D} . Since the Sigmoid function is monotonic, minimizing Eq. (15) implies minimizing the sum of preference term and regularization term. As a result, objective Eq. (15) guides π_θ in a direction that increases the difference between the D-MSE of losing segments and that of the winning segments, while maintaining the overall D-MSE for both winning and losing segments. This is equivalent to increasing the likelihood of winning segments relative to the likelihood of losing segments while holding the likelihood of the whole dataset, which is exactly the goal of alignment phase.

Advantage over Reverse KL Regularization

In fact, Wallace et al. (2023) use the objective Eq. (3), which employs reverse KL regularization to address OOD issues. They provide a different alignment loss as

$$\mathcal{L}_{\text{DPORK}} \approx -\mathbb{E}_{(\sigma^+, \sigma^-) \sim \mathcal{D}_{\text{pref}}} \mathbb{E}_{t \sim \mathcal{U}(1, T), \epsilon_t^+ \sim \mathcal{N}(0, I), \epsilon_t^- \sim \mathcal{N}(0, I)} \text{Sigmoid} \left(-\rho \left(\underbrace{\|\epsilon_t^+ - \epsilon_\theta(\sigma_t^+, t)\|^2 - \|\epsilon_t^- - \epsilon_\theta(\sigma_t^-, t)\|^2}_{\text{preference}} \right) - \underbrace{\|\epsilon_t^+ - \epsilon_{\text{ref}}(\sigma_t^+, t)\|^2 - \|\epsilon_t^- - \epsilon_{\text{ref}}(\sigma_t^-, t)\|^2}_{\text{regularization}} \right) \right). \quad (16)$$

We can see that Eq. (16) contains the same preference term but a different regularization term. This regularization term does not involve trainable policy parameters π_θ . It acts more like introducing a bias within the nonlinear Sigmoid function, rather than directly controlling the D-MSE of π_θ across the entire dataset as Eq. (15). Therefore, Eq. (16) might guide π_θ in an incorrect direction, where the D-MSE of winning segments is significantly lower than that of losing segments, while the overall D-MSE of both winning and losing segments increases. In other words, the likelihood of all segments may decrease, even though the winning segments have a higher likelihood than the losing segments. This leads to the resulting π_θ suffering from OOD issues.

We also provide a toy model to intuitively demonstrate the distinction between reverse KL and forward KL regularization during the alignment process of diffusion policy. As




Figure 2: Samples generated by a toy diffusion model. The left sub-figure displays samples generated by the initial model. The middle and right sub-figures show samples generated by the aligned models with forward and reverse KL regularization, respectively.

shown in Fig. 2, the initial diffusion policy (left) is trained on a dataset following a two-dimensional mixture Gaussian distribution. The centers and radii of blue circles represent the mean and standard deviation of Gaussian components. The reward is defined as the dot product between a sample and the vector $(1, 1)/\sqrt{2}$. We then build preference dataset using the script teacher (Kim et al. 2023), and align the initial policy with the preference dataset using forward KL and reverse KL regularization, respectively. The samples generated by the two aligned models are shown in the middle and right sub-figures of Fig. 2. It is evident that the diffusion model aligned with reverse KL regularization generates a significant number of OOD samples, despite these samples having fairly high reward values. In contrast, the model aligned with forward KL regularization tends to produce samples with relatively high reward values, with most samples not exhibiting OOD issues. In conclusion, forward KL regularization is more effective in avoiding OOD issues during the alignment process of diffusion models.

Related Work

Diffusion Model for Decision Making. Diffusion model is a kind of generative model that exhibits strong capabilities in generating high-dimensional samples (Song, Meng, and Ermon 2020; Lu et al. 2022), e.g., images and text (Kawar et al. 2023; Chen et al. 2023). Diffusion models have been recently adopted in RL to solve various decision-making problems (Zhu et al. 2023). Diffuser (Janer et al. 2022) first proposes to generate state-action trajectories via classifier-guided sampling. Decision-Diffusion (Ajay et al. 2023), MTDiff (He et al. 2023), MetaDiffuser (Ni et al. 2023) and SkillDiffuser (Liang et al. 2024) extend such a framework by using various properties of trajectories for classifier-free guidance, such as task prompts, demonstrations, constraints, skills, representations, and attributes. Other works like Diffusion-QL (Wang, Hunt, and Zhou 2023), EDP (Kang et al. 2023), and IDQL (Hansen-Estruch et al. 2023) treat the conditional diffusion model as a policy and update the policy to maximize the value function. However, the above methods require manually designed reward functions for conditional generation, and learned policy cannot be aligned to preferences. For imitation learning, Diffusion Policy (Chi et al. 2023), VLP (Du et al. 2024), and UniPi (Du et al. 2023) predict future states or

actions based on expert demonstrations. AlignDiff (Dong et al. 2024) tries to align the diffusion model with preferences by first training an attribute strength model and then performing conditional generation for diffusion models. In contrast, we take an alternative way for preference alignment via policy optimization without requiring an attribute model for trajectory labeling. Other works also adopt the diffusion model as a data synthesizer (Lu, Ball, and Parker-Holder 2023), environment model (Yang et al. 2024), and reward generator (Nutti, Franzmeyer, and Henriques 2023) for decision-making, while their contributions are orthogonal to our work.

RLHF and Preference-based RL. RLHF has gained much attention in aligning LLMs with human preference (Ouyang et al. 2022; Bai et al. 2022). RLHF usually contains two stages that learn a reward function first (Bradley and Terry 1952; Lambert et al. 2024) and optimize the LLMs via RL algorithms (Li et al. 2023; Ahmadian et al. 2024). Following this paradigm, Preference-based RL tries to align the policy output with preference data in decision-making problems by first learning the reward function through a BT model (Kim et al. 2023; Lee et al. 2021) and performing policy optimization through RL (Liu et al. 2022). Recently, direct preference learning (DPO) (Rafailov et al. 2024b) and the follow-up works (Rafailov et al. 2024a; Wang et al. 2024) proposed a novel perspective that implicitly optimizes the same objective as RLHF via a KL-constrained reward maximization framework, which forms an implicit reward function. Following this, DPPO (An et al. 2023) and Contrastive Preference Learning (CPL) (Hejna et al. 2023) extend direct preference learning to solve preference-based RL problems via contrastive learning of policy likelihoods or distances. Compared to these methods, we use an advanced diffusion model to enhance the expression ability of policies, while it causes further challenges in a DPO optimization, and we solve it via an approximate objective and additional regularization. The idea of preference alignment has recently been used for diffusion models to improve visual appeal and text alignment in image generation problems through two-stage RLHF (Black et al. 2023; Fan et al. 2023) and DPO (Clark et al. 2023; Wallace et al. 2023; Yang et al. 2023). In contrast, our method focuses on learning diffusion policy through RLHF and proposes the first DPO-based learning framework that learns policies directly from preferences without RL. Regularization methods, particularly reverse KL and forward KL, have been widely discussed in the alignment of LLM (Sun and van der Schaar 2024). Reverse KL and forward KL are known for their mode-seeking and mass-covering characteristics, respectively (Wang et al. 2024). Moreover, forward KL is better suited for balancing alignment performance and generation diversity. In this paper, we highlight another significant advantage of forward KL in the alignment of diffusion policies, namely its enhanced effectiveness in avoiding OOD issues.

Experiments

In this section, we carry out extensive experiments to answer the following three questions. **Q1:** Is FKPD an effective

solver for the PBRL problem? **Q2:** Is forward KL regularization crucial to FKPD, and is it more effective than reverse KL regularization? **Q3:** How does FKPD perform under different hyperparameter settings? (Please refer to Appendix 2.1)

We evaluate the performance of FKPD on two well-known benchmarks: MetaWorld robotics tasks (Yu et al. 2020) and D4RL locomotion (Fu et al. 2020) tasks. The first benchmark uses the same offline dataset as Contrastive Preference Learning (CPL) (Hejna et al. 2023). This dataset is collected on six tasks from the simulated MetaWorld robotics environment (Yu et al. 2020). For each task, the preference-free dataset is collected by rolling out 2500 episodes of length 250 with a sub-optimal stochastic policy. The preference dataset is collected by uniformly sampling segments of length 64 from the preference-free dataset and adding preference labels using script teacher (Kim et al. 2023). In terms of the D4RL benchmark, we generate the preference dataset from the original dataset using the same procedure as (Hejna et al. 2023). For further details about the dataset, please refer to Appendix 3.1.

For baseline methods, we choose three strong baselines. The first baseline is supervised fine-tuning (**SFT**) (Lee et al. 2018), where a policy is first trained with behavior clone on all segments in $\mathcal{D}_{\text{pref}}$, then further fine-tuned on only the winning segments in $\mathcal{D}_{\text{pref}}$. The second baseline is Preference-IQL (**P-IQL**) (Kim et al. 2023), which learns a preference model from $\mathcal{D}_{\text{pref}}$, then learns a policy to maximize it with Implicit Q -learning (IQL) (Kostrikov, Nair, and Levine 2022), an efficient offline RL algorithm. The third one is **CPL**, a state-of-the-art baseline that also uses DPO for the alignment phase while employs only a basic Gaussian policy.

		SFT	P-IQL	CPL	FKPD
2.5k Dense	BinPicking	66.9±2.1	70.6±4.1	80.0±2.5	77.0±3.6
	ButtonPress	21.6±1.6	16.2±5.4	24.5±2.1	34.7±2.1
	DoorOpen	63.3±1.9	69.0±6.2	80.0±6.8	87.7±1.2
	DrawerOpen	62.6±2.4	71.1±2.3	83.6±1.6	90.0±2.7
	PlateSlide	41.6±3.5	49.6±3.4	61.1±3.0	58.3±1.2
	SweepInto	51.9±2.1	60.6±3.6	70.4±3.0	76.5±3.5
20k Sparse	BinPicking	67.0±4.9	75.0±3.3	83.2±3.5	76.3±1.5
	ButtonPress	21.4±2.7	19.5±1.8	29.8±1.8	33.3±4.5
	DoorOpen	63.6±2.4	79.0±6.6	77.9±5.0	84.3±2.1
	DrawerOpen	63.5±0.9	76.2±2.8	79.1±5.0	91.3±2.3
	PlateSlide	41.9±3.1	55.5±4.2	56.4±3.9	66.3±7.4
	SweepInto	50.9±3.2	73.4±4.2	81.2±1.6	76.3±2.1

Table 1: Success rates (in percent) of all methods for six MetaWorld tasks on different datasets. Here 2.5k Dense and 20k sparse represent two different preference datasets, with detailed explanations provided in Appendix 3.1.

Performance of FKPD

The results of FKPD and competitors for MetaWorld are presented in Table 1. The architecture of the policy network, along with evaluation details and hyperparameter settings,

	SFT	P-IQL	CPL	FKPD
expert	walk2d	82.9±8.5	109.8±0.4	107.9±0.2
	hopper	48.3±5.0	84.5±4.1	80.5±3.4
	halfche	68.5±4.7	83.6±3.8	94.7±0.3
replay	walk2d	33.9±7.1	71.2±10.3	48.3±3.7
	hopper	57.6±5.6	68.9±33.8	111.2±0.2
	halfche	34.8±1.8	42.3±0.5	45.3±0.1

Table 2: Average norm score of all methods for D4RL locomotion tasks on different datasets.

are provided in Appendix 3-4. When we use the dense comparison data, FKPD achieves the best performance in 4 of 6 tasks. Especially in *Button Press*, *Door Open*, and *Drawer Open*, FKPD has significant advantages over its competitors. When using sparser comparison data, FKPD also outperforms its competitors in 4 of 6 tasks, with a substantial margin in *Button Press*, *Door open* and *Drawer Open*. In conclusion, FKPD is the method with the best overall performance in this dataset. Compared to SFT, FKPD makes more effective use of the information in $\mathcal{D}_{\text{pref}}$, leading to better performance. Compared with P-IQL, FKPD directly optimizes the policy using preference data, without learning an explicit preference score function. This eliminates the risk associated with uncertain score functions. Compared with CPL, FKPD adopts a more expressive diffusion policy, making it more suitable for complex manipulation tasks.

The results of FKPD and competitors for D4RL are demonstrated in Table 2. For this benchmark, we also use a critic as (Wang, Hunt, and Zhou 2023), which is directly learned from the preference dataset (Hejna and Sadigh 2024) without the need of a reward model, to filter out sub-optimal candidate actions. We observe that the proposed FKPD outperforms CPL in 5 out of 6 D4RL locomotion tasks, further demonstrating the advantages of employing a diffusion strategy. Additionally, it achieves competitive performance with the two-stage P-IQL without the need of training an explicit reward model. In a word, FKPD is an effective solver for D4RL locomotion tasks.

Ablation of Regularization

We next provide a detailed comparison of two aforementioned regularization methods: forward KL and reverse KL regularization. To this end, we compare FKPD and Reverse KL regularized for aligning Diffusion Policies (RKPD) (Wallace et al. 2023) on three representative MetaWorld (20k sparse) tasks, namely, *Door Open*, *Drawer Open* and *Button Press*; as well as three representative D4RL locomotion tasks, namely, *Walker2d-medium-replay*, *hopper-medium-replay* and *halfcheetah-medium-replay*. Moreover, we also provide the results of DPO without regularization (referred to as NRPD) for reference. Here we use a novel metric to evaluate the performance of the alignment phase for these three methods. For each method, we record the policy’s average success rate (or average return for D4RL tasks) on a task as U_0 at the beginning of the alignment phase. After the alignment phase, we evaluate its average success rate

	FKPD	RKPD	NRPD
Door Open	40.0±3.5	21.7±8.8	-42.1±25.6
Drawer Open	47.2±3.7	6.9±3.4	-73.7±24.5
Button Press	58.6±21.6	60.9±21.0	-100.0±0.0
Walk2d-med-rep	75.4±5.0	-88.7±11.9	-95.1±6.8
hopper-med-rep	64.8±12.5	-33.9±17.9	-98.4±1.0
halfche-med-rep	14.9±1.2	-45.3±54.3	-100.6±9.4

Table 3: Average improvement factor (in percentage) of three DPO methods on six MetaWorld and D4RL tasks.

(or average return) again, recorded as U_1 . Then, we define the improvement factor as $F_{\text{im}} = (U_1 - U_0)/U_0$. All results are present in Table 3.

For three MetaWorld tasks, we observe that NRPD exhibits the worst performance among the three tasks, especially in *Button Press*, where it completely fails. This demonstrates the importance of regularization terms in diffusion policy learning. In addition, FKPD outperforms RKPD in the *Door Open* and *Drawer Open* tasks, performs slightly worse in the *Button Press* task, and overall demonstrates better performance. For three D4RL tasks, We observe that FKPD demonstrates a more pronounced advantage. At this point, RKPD and NRPD can not even enhance policy performance during the alignment phase. This implies that regularization is of greater importance in D4RL tasks.

Next we conduct an in-depth analysis of why forward KL regularization outperforms reverse KL regularization for the aligning of diffusion policy. To this end, we track three variables in the process of alignment. The first two variables are average D-MSE of winning segments and losing segments, abbreviated as E_{winning} and E_{losing} , respectively.

They are capable of measuring the variation in the likelihood of winning segments and losing segments during the alignment process of diffusion policy. The third variable is implicit accuracy, which is written as

$$I_{\text{acc}} = \frac{1}{B} \sum_{\sigma^+ \sigma^- \in \mathcal{B}_{\text{pref}}} \mathbb{I} \left(\left\| \epsilon_t^+ - \epsilon_\theta(\sigma_t^+, t) \right\|^2 - \left\| \epsilon_t^- - \epsilon_\theta(\sigma_t^-, t) \right\|^2 < 0 \right), \quad (17)$$

where \mathbb{I} is the indicator function, $\mathcal{B}_{\text{pref}}$ is a batch of segments with preference labels, B is the size of $\mathcal{B}_{\text{pref}}$. I_{acc} represents the estimated probability that the likelihood of π_θ generating winning segments is greater than that of generating losing segments. Fig. 3 illustrates these three variables for FKPD, RKPD, and NRPD in *Drawer Open* and *Walker2d-medium-replay*. We also provide these variables of the pre-trained policy from the behavior clone phase for reference. We start with NRPD. Although it ensures that the likelihood of generating winning segments increases relative to the likelihood of generating losing segments (as indicated by I_{acc} values) during training, the E_{winning} and E_{losing} values reveal that the likelihood of both positive and negative segments decreases significantly. This suggests that the NRPD method suffers from a severe OOD issue. FKPD, while ensuring an

increased likelihood of generating positive segments relative to negative segments, also effectively maintains the values of E_{winning} and E_{losing} without significant increases. This allows it to preserve the overall likelihood of both positive and negative segments, thereby addressing the OOD issue present in NRPD. RKPd, similar to NRPd, significantly increases the likelihood of positive segments relative to negative segments. However, the likelihood for both the positive and negative segments decreases noticeably compared to the reference values. Thus, it fails to address the OOD challenge as effectively as the FKPD.

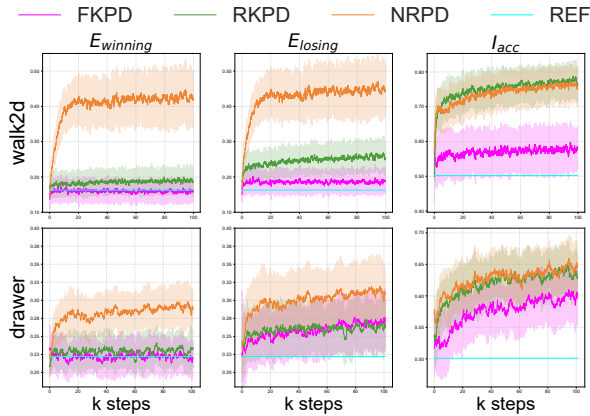


Figure 3: Average D-MSE and implicit accuracy for FKPD, NRPD, and RKPd during the alignment phase of *Drawer Open* (bottom row) and *Walker2d-medium-replay* (top row). We also provide the same variables of the pretrained policy on the preference-free dataset \mathcal{D} for reference.

Conclusions

In this paper, we propose FKPD, a method to align a diffusion policy with preference data. To tackle the OOD issue during the alignment phase, we employ the forward KL regularization, which demonstrates superior performance compared to the reverse KL regularization in the experiments. FKPD achieves the state-of-the-art performance on open MetaWorld and D4RL locomotion tasks, indicating that it is an efficient preference aligning method. Currently, we have only validated the effectiveness of FKPD for single-modal observations, that is, the joint states of embodiment, and single-task scenarios. In future work, we aim to extend FKPD to encompass multi-task and multi-modal data. Moreover, other efficient regularization methods, such as those proposed in (Huang et al. 2024; Cen et al. 2024), have been introduced for aligning LLMs. In our future work, we plan to analyze the performance of these new regularization methods for aligning diffusion policies.

Acknowledgments

This study is supported by National Natural Science Foundation of China (Grant No.62306242). Zhao extends gratitude to professor Chenjia for providing in-depth guidance

on this work. Thanks are also extended to Chenyou and the other authors for their outstanding contributions to the research.

References

- Ahmadian, A.; Cremer, C.; Gallé, M.; Fadaee, M.; Kreutzer, J.; Üstün, A.; and Hooker, S. 2024. Back to Basics: Revisiting REINFORCE Style Optimization for Learning from Human Feedback in LLMs. *arXiv preprint arXiv:2402.14740*.
- Ajay, A.; Du, Y.; Gupta, A.; Tenenbaum, J. B.; Jaakkola, T. S.; and Agrawal, P. 2023. Is Conditional Generative Modeling all you need for Decision Making? In *The Eleventh International Conference on Learning Representations*.
- An, G.; Lee, J.; Zuo, X.; Kosaka, N.; Kim, K.-M.; and Song, H. O. 2023. Direct Preference-based Policy Optimization without Reward Modeling. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Bai, Y.; Jones, A.; Ndousse, K.; Askell, A.; Chen, A.; Das-Sarma, N.; Drain, D.; Fort, S.; Ganguli, D.; Henighan, T.; et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Black, K.; Janner, M.; Du, Y.; Kostrikov, I.; and Levine, S. 2023. Training Diffusion Models with Reinforcement Learning. In *ICML 2023 Workshop on Structured Probabilistic Inference & Generative Modeling*.
- Bradley, R. A.; and Terry, M. E. 1952. Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika*, 39(3/4): 324–345.
- Cen, S.; Mei, J.; Goshvadi, K.; Dai, H.; Yang, T.; Yang, S.; Schuurmans, D.; Chi, Y.; and Dai, B. 2024. Value-Incentivized Preference Optimization: A Unified Approach to Online and Offline RLHF. *arXiv preprint arXiv:2405.19320*.
- Chen, J.; Huang, Y.; Lv, T.; Cui, L.; Chen, Q.; and Wei, F. 2023. TextDiffuser: Diffusion Models as Text Painters. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Chi, C.; Feng, S.; Du, Y.; Xu, Z.; Cousineau, E.; Burchfiel, B.; and Song, S. 2023. Diffusion Policy: Visuomotor Policy Learning via Action Diffusion. In *Robotics: Science and Systems*.
- Clark, K.; Vicol, P.; Swersky, K.; and Fleet, D. J. 2023. Directly fine-tuning diffusion models on differentiable rewards. *arXiv preprint arXiv:2309.17400*.
- Dong, Z.; Yuan, Y.; HAO, J.; Ni, F.; Mu, Y.; ZHENG, Y.; Hu, Y.; Lv, T.; Fan, C.; and Hu, Z. 2024. AlignDiff: Aligning Diverse Human Preferences via Behavior-Customisable Diffusion Model. In *The Twelfth International Conference on Learning Representations*.
- Du, Y.; Yang, S.; Dai, B.; Dai, H.; Nachum, O.; Tenenbaum, J. B.; Schuurmans, D.; and Abbeel, P. 2023. Learning Universal Policies via Text-Guided Video Generation. In *Thirty-seventh Conference on Neural Information Processing Systems*.

- Du, Y.; Yang, S.; Florence, P.; Xia, F.; Wahid, A.; brian ichter; Sermanet, P.; Yu, T.; Abbeel, P.; Tenenbaum, J. B.; Kaelbling, L. P.; Zeng, A.; and Tompson, J. 2024. Video Language Planning. In *The Twelfth International Conference on Learning Representations*.
- Fan, C.; Bai, C.; Shan, Z.; He, H.; Zhang, Y.; and Wang, Z. 2024. Task-agnostic Pre-training and Task-guided Fine-tuning for Versatile Diffusion Planner. *arXiv preprint arXiv:2409.19949*.
- Fan, Y.; Watkins, O.; Du, Y.; Liu, H.; Ryu, M.; Boutilier, C.; Abbeel, P.; Ghavamzadeh, M.; Lee, K.; and Lee, K. 2023. Reinforcement Learning for Fine-tuning Text-to-Image Diffusion Models. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Fu, J.; Kumar, A.; Nachum, O.; Tucker, G.; and Levine, S. 2020. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*.
- Fujimoto, S.; Meger, D.; and Precup, D. 2019. Off-policy deep reinforcement learning without exploration. In *International conference on machine learning*, 2052–2062. PMLR.
- Gu, J.; Xiang, F.; Li, X.; Ling, Z.; Liu, X.; Mu, T.; Tang, Y.; Tao, S.; Wei, X.; Yao, Y.; Yuan, X.; Xie, P.; Huang, Z.; Chen, R.; and Su, H. 2023. ManiSkill2: A Unified Benchmark for Generalizable Manipulation Skills. In *The Eleventh International Conference on Learning Representations*.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, 1861–1870. PMLR.
- Hansen-Estruch, P.; Kostrikov, I.; Janner, M.; Kuba, J. G.; and Levine, S. 2023. Idql: Implicit q-learning as an actor-critic method with diffusion policies. *arXiv preprint arXiv:2304.10573*.
- He, H.; Bai, C.; Pan, L.; Zhang, W.; Zhao, B.; and Li, X. 2024. Learning an actionable discrete diffusion policy via large-scale actionless video pre-training. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- He, H.; Bai, C.; Xu, K.; Yang, Z.; Zhang, W.; Wang, D.; Zhao, B.; and Li, X. 2023. Diffusion Model is an Effective Planner and Data Synthesizer for Multi-Task Reinforcement Learning. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Hejna, J.; Rafailov, R.; Sikchi, H.; Finn, C.; Niekum, S.; Knox, W. B.; and Sadigh, D. 2023. Contrastive Preference Learning: Learning from Human Feedback without RL. *arXiv preprint arXiv:2310.13639*.
- Hejna, J.; and Sadigh, D. 2024. Inverse preference learning: Preference-based rl without a reward function. *Advances in Neural Information Processing Systems*, 36.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Ho, J.; and Salimans, T. 2021. Classifier-Free Diffusion Guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*.
- Huang, A.; Zhan, W.; Xie, T.; Lee, J. D.; Sun, W.; Krishnamurthy, A.; and Foster, D. J. 2024. Correcting the Mythos of KL-Regularization: Direct Alignment without Overoptimization via Chi-Squared Preference Optimization. *arXiv preprint arXiv:2407.13399*.
- Janner, M.; Du, Y.; Tenenbaum, J. B.; and Levine, S. 2022. Planning with Diffusion for Flexible Behavior Synthesis. In *ICML*, volume 162 of *Proceedings of Machine Learning Research*, 9902–9915. PMLR.
- Kang, B.; Ma, X.; Du, C.; Pang, T.; and Yan, S. 2023. Efficient diffusion policies for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 36.
- Kawar, B.; Zada, S.; Lang, O.; Tov, O.; Chang, H.; Dekel, T.; Mosseri, I.; and Irani, M. 2023. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6007–6017.
- Kim, C.; Park, J.; Shin, J.; Lee, H.; Abbeel, P.; and Lee, K. 2023. Preference Transformer: Modeling Human Preferences using Transformers for RL. In *The Eleventh International Conference on Learning Representations*.
- Knox, W. B.; Hatgis-Kessell, S.; Booth, S.; Niekum, S.; Stone, P.; and Allievi, A. G. 2024. Models of human preference for learning reward functions. *Transactions on Machine Learning Research*.
- Kostrikov, I.; Nair, A.; and Levine, S. 2022. Offline Reinforcement Learning with Implicit Q-Learning. In *International Conference on Learning Representations*.
- Lambert, N.; Pyatkin, V.; Morrison, J.; Miranda, L.; Lin, B. Y.; Chandu, K.; Dziri, N.; Kumar, S.; Zick, T.; Choi, Y.; et al. 2024. Rewardbench: Evaluating reward models for language modeling. *arXiv preprint arXiv:2403.13787*.
- Lange, S.; Gabel, T.; and Riedmiller, M. 2012. Batch reinforcement learning. In *Reinforcement learning: State-of-the-art*, 45–73. Springer.
- Lee, C.; Panda, P.; Srinivasan, G.; and Roy, K. 2018. Training deep spiking convolutional neural networks with stdp-based unsupervised pre-training followed by supervised fine-tuning. *Frontiers in neuroscience*, 12: 373945.
- Lee, K.; Smith, L.; Dragan, A.; and Abbeel, P. 2021. B-Pref: Benchmarking Preference-Based Reinforcement Learning. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- Li, Z.; Xu, T.; Zhang, Y.; Yu, Y.; Sun, R.; and Luo, Z.-Q. 2023. ReMax: A Simple, Effective, and Efficient Reinforcement Learning Method for Aligning Large Language Models. *arXiv e-prints*, arXiv–2310.
- Liang, Z.; Mu, Y.; Ma, H.; Tomizuka, M.; Ding, M.; and Luo, P. 2024. Skilldiffuser: Interpretable hierarchical planning via skill abstractions in diffusion-based task execution. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Liu, R.; Bai, F.; Du, Y.; and Yang, Y. 2022. Meta-Reward-Net: Implicitly Differentiable Reward Learning for Preference-based Reinforcement Learning. In *Advances in Neural Information Processing Systems*, volume 35, 22270–22284.

- Lu, C.; Ball, P. J.; and Parker-Holder, J. 2023. Synthetic Experience Replay. In *Workshop on Reincarnating Reinforcement Learning at ICLR 2023*.
- Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2022. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35: 5775–5787.
- Ni, F.; Hao, J.; Mu, Y.; Yuan, Y.; Zheng, Y.; Wang, B.; and Liang, Z. 2023. MetaDiffuser: Diffusion Model as Conditional Planner for Offline Meta-RL. In *ICML*, volume 202 of *Proceedings of Machine Learning Research*, 26087–26105. PMLR.
- Nichol, A. Q.; and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, 8162–8171. PMLR.
- Nuti, F. P. C.; Franzmeyer, T.; and Henriques, J. F. 2023. Extracting Reward Functions from Diffusion Models. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- OpenAI. 2023. GPT-4 Technical Report.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- Pearce, T.; Rashid, T.; Kanervisto, A.; Bignell, D.; Sun, M.; Georgescu, R.; Macua, S. V.; Tan, S. Z.; Momennejad, I.; Hofmann, K.; and Devlin, S. 2023. Imitating Human Behaviour with Diffusion Models. In *The Eleventh International Conference on Learning Representations*.
- Qin, R.-J.; Zhang, X.; Gao, S.; Chen, X.-H.; Li, Z.; Zhang, W.; and Yu, Y. 2022. NeoRL: A Near Real-World Benchmark for Offline Reinforcement Learning. In *Neural Information Processing Systems Datasets and Benchmarks Track*.
- Rafailov, R.; Hejna, J.; Park, R.; and Finn, C. 2024a. From r to Q^* : Your Language Model is Secretly a Q-Function. *arXiv preprint arXiv:2404.12358*.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2024b. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Reuss, M.; Li, M.; Jia, X.; and Lioutikov, R. 2023. Goal-conditioned imitation learning using score-based diffusion policies. In *Robotics: Science and Systems*.
- Shin, D.; Dragan, A. D.; and Brown, D. S. 2023. Benchmarks and Algorithms for Offline Preference-Based Reward Learning. *Trans. Mach. Learn. Res.*, 2023.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising Diffusion Implicit Models. In *International Conference on Learning Representations*.
- Sun, H.; and van der Schaar, M. 2024. Inverse-RLignment: Inverse Reinforcement Learning from Demonstrations for LLM Alignment. *arXiv preprint arXiv:2405.15624*.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Wallace, B.; Dang, M.; Rafailov, R.; Zhou, L.; Lou, A.; Pushwalkam, S.; Ermon, S.; Xiong, C.; Joty, S.; and Naik, N. 2023. Diffusion model alignment using direct preference optimization. *arXiv preprint arXiv:2311.12908*.
- Wang, C.; Jiang, Y.; Yang, C.; Liu, H.; and Chen, Y. 2024. Beyond Reverse KL: Generalizing Direct Preference Optimization with Diverse Divergence Constraints. In *The Twelfth International Conference on Learning Representations*.
- Wang, Z.; Hunt, J. J.; and Zhou, M. 2023. Diffusion Policies as an Expressive Policy Class for Offline Reinforcement Learning. In *The Eleventh International Conference on Learning Representations*.
- Yang, K.; Tao, J.; Lyu, J.; Ge, C.; Chen, J.; Li, Q.; Shen, W.; Zhu, X.; and Li, X. 2023. Using Human Feedback to Fine-tune Diffusion Models without Any Reward Model. *arXiv preprint arXiv:2311.13231*.
- Yang, S.; Du, Y.; Ghasemipour, S. K. S.; Tompson, J.; Kaelbling, L. P.; Schuurmans, D.; and Abbeel, P. 2024. Learning Interactive Real-World Simulators. In *The Twelfth International Conference on Learning Representations*.
- Yu, T.; Quillen, D.; He, Z.; Julian, R.; Hausman, K.; Finn, C.; and Levine, S. 2020. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on robot learning*, 1094–1100. PMLR.
- Yu, X.; Bai, C.; He, H.; Wang, C.; and Li, X. 2024. Regularized Conditional Diffusion Model for Multi-Task Preference Alignment. *arXiv preprint arXiv:2404.04920*.
- Yuan, H.; Huang, K.; Ni, C.; Chen, M.; and Wang, M. 2023. Reward-directed conditional diffusion: Provable distribution estimation and reward improvement. *Advances in Neural Information Processing Systems*, 36.
- Yuan, X.; Shang, Z.; Wang, Z.; Wang, C.; Shan, Z.; Qi, Z.; Zhu, M.; Bai, C.; and Li, X. 2024a. Preference Aligned Diffusion Planner for Quadrupedal Locomotion Control. *arXiv preprint arXiv:2410.13586*.
- Yuan, Y.; HAO, J.; Ma, Y.; Dong, Z.; Liang, H.; Liu, J.; Feng, Z.; Zhao, K.; and ZHENG, Y. 2024b. Uni-RLHF: Universal Platform and Benchmark Suite for Reinforcement Learning with Diverse Human Feedback. In *The Twelfth International Conference on Learning Representations*.
- Ze, Y.; Zhang, G.; Zhang, K.; Hu, C.; Wang, M.; and Xu, H. 2024. 3d diffusion policy. In *Robotics: Science and Systems*.
- Zhu, Z.; Zhao, H.; He, H.; Zhong, Y.; Zhang, S.; Yu, Y.; and Zhang, W. 2023. Diffusion models for reinforcement learning: A survey. *arXiv preprint arXiv:2311.01223*.